The Real Options Approach to Standardization

Mark Gaynor
Harvard University
gaynor@eecs.harvard.edu

Scott Bradner
Harvard University
sob@eecs.harvard.edu

ABSTRACT

In this paper, we propose a new model of technology standardization under market uncertainty and show how its value is quantifiable using the theory of real options. Our options-based approach to standardization shows that a rational way to standardize some IT technology in uncertain markets is with correct structure and proper staging of the standard. First, highly modularized standards provide a higher option value because of the ability to pick and choose the best modules to change at a fine granularity. Secondly, a modular structure that promotes easy and non-disruptive parallel experimentation (such as end-2-end applications) enhances the option value by providing a larger field of options from which to select. Lastly, allowing the standard to evolve along with the customers' expectations of the technology is a good strategy to match standards with uncertain user markets.

1. Introduction

Product development involving computers, software, and networking has had a profound impact on theories of innovation and product development. Technology in these products changes very rapidly when compared to that in many traditional industries such as automobiles [1], or production of Television picture tubes [2]. Just imagine autos that double in speed every 18 months, similar to the performance increase in microprocessors. The faster evolution of these technologies does not fit traditional product development theories that depend on periods of disruptive innovation followed by less drastic incremental changes [3,4,5,6], given that the period some technologies such as computers and information are stable is very short. A new breed of models [7,8,9] views the evolution of technology as a continuum of changes, not the punctuated equilibrium of the past.

Customer expectations co-evolve with technology change at today's faster pace, creating uncertainty in consumer preferences. Clark [1] points out that when a new technology is born, customers have no education about the technology and tend to view it in the context of what is being replaced. The evolution of customer expectations of the web is a good example. At first, the web was mainly a tool for researchers sharing information, the important service attribute being that the data existed, and is accessible by heterogeneous computer systems. Only later, as the interactive nature of the web matured, did consumers become more sophisticated in the services they demanded. Now information layout, e-commerce, and usability have become important attributes of web-based services.

Changes in product development in the computer age are parallel to alterations in effective standardization of technology. With slower-moving technologies, standardization occurs in the relatively stable period after technology selection. But this stable period is short, or non-existent in fast changing technologies like DRAMs, where useful standards must be timely, and produced in a few months, not years [10]. The uncertainty created by evolving customer preferences means that (once created) standards must have the ability to evolve along with the end users of the standard.

Recently, standards have become more important to business, thus causing strategic management of companies' standard's policies to play an increasingly important role in formulation of overall corporate strategy. Evidence of this is the increased memberships of fee-based industry consortiums and alliances such as W3C, ATM forum, and X/Open. A report to the chairman of W3C in 1999 shows an increase from 30 to 370 members. Cargill [11] also notes that at one point in the 1990's, consortiums for business-based standards increased by two per month. This shows that the demand for standards outstrips the supply produced by traditional SDOs such as ISO and ANSI.

E-commerce was the big buzz of the late 90's with companies like Amazon.com, ebay, and others reaching market capitalization far beyond expectations. In order for e-commerce to function, standards must be in place. These include: web and networking standards, security, as well as standards like XML that structure the data exchanged between vendor and customer (or other vendor). Without standards, e-commerce is not possible.

Networking, IT, and other technology standards in areas of uncertain user preferences are not static documents, but dynamic complex adaptive systems (CAS) that must interact and change within their environment. Factors causing these standards to behave as CAS are: increased number of users with diverse, fast-changing, unpredictable requirements, and uncertainty of implementation of a standard. For better success in this new uncertain environment, standards must follow principles similar to how evolution and natural selection picks the fittest organisms, but, with the market[1] as the selector picking the fittest technology in terms of the users (or users and vendors). This paradigm is best suited to describe the standardization process in today's ever-changing dynamic environment.

A standard development methodology that promotes a broad range of experimentation combined with market selection will better serve end users by involving them in the standardization process. Promoting experimentation with new proposed standards and standardizing the technology adopted by most of the community decreases the risk of an unaccepted standard. Design principles such as the end-2-end argument that push intelligence to the network's edge help promote this type of experimentation because they broaden the range of participants able to innovate.

The theory of real options applied to standards gives the right, but not the obligation, to follow a path of standardization. For example, the IETF category, "proposed standard" is a profile of options; it allows the Internet community to choose the standards that succeed by exercising the option to implement the standard and provide the services enabled by the standard. The standardization option standardizes what the market selects.

In this paper, we propose a new prescriptive model of technology standardization under uncertainty and show how its value is quantifiable using the theory of real options, a proven methodology for management of non-financial assets under uncertainty. Our model is simple and intuitive: start with simple standards structured in a modular layered architecture, then let them evolve, with the market acting as the selection mechanism. Our model of standardization shows how modularity creates value from uncertainty by maximizing the choice of options along the standardization path. We explain how to apply this framework to the development of communication protocol standards, but do not provide a numerical example. We examine several different levels in the hierarchy of standards. First, we examine the architecture of a protocol stack to show the value of a layered modular structure. This value is realized in terms of both choosing the best protocols for inclusion into a standard suite and in terms of providing an environment that allows building the best implementation of the standard. Our model shows that modularity (up to a point) intrinsically creates greater value than an interconnected design, with the number and diversity of services depending on the particular layer of the stack. We argue that a "thin/thick" structure (see Section 4.1) where the network only provides basic transport services pushes services and applications to intelligent end points creating an environment conducive to experimentation. For example, the Network layer (IP) should contain the fewest protocols that provide only the most basic services, while the application-layer should contain the most protocols with the most diversity in terms of services offered. Next, we discuss the value created by applying the methodology of introducing simple protocol suites (and protocols) and evolving the stack by creating new protocols or altering existing ones. Our theory shows that the evolutionary approach to development of entire protocol stacks, and protocols within each layer, maximizes the expected value of the standard.

This paper should be of interest to both academics, and practitioners interested in standards, or the architecture of protocols. For the academic, this paper presents a new idea: using the theory of real options to value a standard. This paper is the tip of the iceberg, there is much further research in this area. For those who create standards, this paper presents a new mindset - think in terms of keeping your options open, the more options you allow, the better your expected outcome.

The structure of this paper is as follows. In section 2, we place standards in the context of complex adaptive systems. Section 3 is the methodology showing how this theory is based on previous work about how markets select standards, and how the theory of real options has been used to show the value of modularity in computer systems design when the technology outcome of the system components is uncertainty. Next, in Section 4, our model is discussed and the theory is explained showing how the theory of real options helps quantify the value of modularity and evolution in standards for complex fast changing technology. Last, in Section 5, general rules of protocol standardization are discussed and generalized to other standards.

2. IT Standards: Complex Adaptive Systems in an Uncertain World

Modern technology is complex with much uncertainty; users and vendors each have different needs for the standardization process and competition exists both between and within standards. The regulatory environment is now different, allowing support for a more open standardization process, and providing incentives for those creating the standards. Our theory depends on competing technical solutions for standards to provide the users with options. Unpredictable and dynamic user needs cause vendors to have incomplete knowledge of how a particular standard will mature and be used. While general services are predictable in many instances, the particular feature set, and implementation is often not. Email is a good example

[1] This market may be users, or a group of interested parties (such as the IETF), as discussed in Section 3.1.

of this, the demand was clear, but it took several generations of competing service offerings to converge to a standard based solution. As noted by Clark [1], customers do not have the knowledge of with new technologies to understand the possibilities. These customer expectations must evolve along with the technology; the interaction between the technology and consumer preferences is very complex. Similar to new views of product development [8, 12, 7], standards fit into the context of Complex Adaptive Systems (CAS) [13, 14, 15]. This implies effective standards must evolve and have a selection process to pick from many competing options.

Uncertainty about the success of a proposed standard is one reason that standards need to start out simple, but display the flexibility (the cost of this flexibility is discussed in Section 4.4.1) to evolve within a continuously changing environment. For our model, we only examine uncertainty (and the associated risk) in market prediction and do not address economic or technology uncertainty.

2.1 Uncertainty

The market for "standards-based products" can be very dynamic and hard to predict. We are interested in need-driven standards where users' needs are a moving target, making even accurate short-term predictions difficult. It may seem a contradiction to claim that demand-pull standards can have elements of unpredictability, but consider email, clearly a predicted success, but the standard that became popular was not the predicted X.400 suite, but the Internet scheme. Many services have clear demand, but uncertainty in the particular feature set and implementation that will work best for most users. Furthermore, firms sometimes get it wrong even in a well-defined market. OSI is a good example of this, the market existed, and demands were clear: interoperability between heterogeneous networks and computer systems. OSI transport, the suite of communications protocols developed by the ISO and championed by all the major vendors and governments (including ours), is dead. The ISO failed to produce a standard accepted by users. Market miss-prediction incurs the risk of introduction of a proprietary solution to meet the market demand.

There are many examples of how vendors are unable to predict what will happen in today's world. Nobody guessed the WWW (based on standards) would be the "killer application" that popularized the Internet, or the dramatic impact the WWW is having on society. The success of the entire Internet and the value created by it vastly surpasses any estimates its creators could imagine (even in their wildest dreams). Technologies like ISDN, SMDS and ATM did not meet the predictions of the experts. These examples show the complexity of the standardization environment for networking systems. Thus far it has been hard to predict which standards become successful and which ones fail.

Even successful standards mature in unforeseen ways. Frame relay is successful in low bandwidth application, but was developed as a medium and high-speed WAN service. ATM failed to reach the desktop (in terms of ATM packets reaching the PC) as expected, but has become a viable solution within the core fabric of high-speed IP routers, and recently in providing DSL. It is precisely this unpredictability of which standards succeed and which applications use what standards that requires a new paradigm.

3. Methodology

This work is theoretical and draws from two main areas of research. First, work by Vercoulen [16] discusses modularity in standards, and how to select standards in dynamic complex industries. Next, research derived from the theory of options shows the value of modularity in computer systems. While based on theory, the empirical evidence supports our model (at a high level of analysis - see Section 4.6). The successful networking protocols have evolved in a modular structure, but the protocols that became popular were not aways those predicted by the industry pundits.

3.1 Modularity and Selection of Standards

Vercoulen [16] discusses how modularity in standards adds value by creating standards that are complex, but able to react to dynamic change quickly. It discusses how modular standards may work best with a combination of market and negotiated selection. Negotiations sometimes help develop complex standards that fit together, but this negotiation process may be too lengthy for dynamic markets. By combining both selection modes, complex working standards can be created in a timely manner. This work classifies complex modular standards as complex dynamic systems and builds a base for our theory.

Modularity of complex standards have advantages and disadvantages. Vercouolen points out benefits of modularity such as: modularity allows specialization where different parties develop different modules, scalability of the system, and innovation by including new modules. However, also discussed are the negative aspect of modular systems such as: coordination failures between modules, resources required to link and coordinate modules add to system overhead, and connecting modules into a cost effective system is non-trivial. Our work focuses on the advantage modularity gives to innovation, while accounting for the additional expense of the modularity as discussed in Section 4.4.1.

3.2 Theory of Options

The theory of options has proven useful for managing financial risk in uncertain environments. To see how options can limit risk, consider the classic call option: it gives the right, but not the obligation, to buy a security at a fixed date in the future, with the price determined in the past. Buying a call option is the equivalent of betting that the underlying security will rise in value more than the price of acquiring the option. The option limits the downside risk, but not the upside gain, thus providing a non-linear payback, unlike owning the security. This implies that options provide increasing value as the uncertainty of the investment grows (i.e. as variance in the distribution describing the value of the security increases), since the downside risk is capped without limiting the upside potential.

Figure 1 shows graphically how this works. The non-linear payback of the option is the solid line, while the linear pay out of owning the stock is the dashed line. The option holder is able to look at the price of the security when the option is due and decide whether to exercise the option to buy the stock. It is the historical variability of the stock price, not the security price that determines the value of the option. This protects the option holder by limiting the loss to the cost of acquiring the option no matter how low the stock price falls. Some risk-adverse investors prefer this type of non-linear payback that caps the downside risk, but leaves the upside gain unaltered.

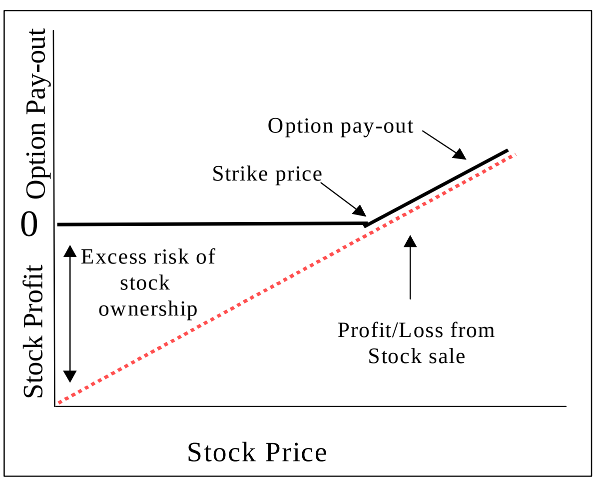

Figure 1 **Option Pay-out**

This theory of options is extendable to options on real (non-financial) assets [17]. Real options provide a structure linking strategic planning and financial strategy. Similar to financial options, real options limit the downside risk of an investment decision without limiting the upside potential. In many cases, this approach shows a greater potential expected value than the standard discounted cash flow analysis performed in most corporate environments. This theory is useful in examining a plethora of situations in the real world such as staged investment in IT infrastructure [18], oil field expansion, developing a drug [17], and even showing the value of modularity in designing computer systems [12].

Staging the investment required to build large IT/Telecommunications systems provides an option at each stage of the investment. This option is whether to continue the investment or not, and is based on the most current information available about the uncertain market, economy, and technical attributes of the project. Starting out small, and evolving the project at various stages allows making more focused and relevant decisions which in turn increase the expected value of a staged implementation over that of the single stage scheme.

In "Design Rules", Baldwin and Clark apply this theory to study modularization in the computer industry. They show how modularization of computer systems design (like the IBM 360) has tremendously changed the industry. A modularly designed computer consists of components that have defined interfaces. Because each component conforms to its interface rules, modules that follow the defined interface are interchangeable. In contrast, an interconnected system has no swappable components because only a single massive component exists.

To see how a modular design provides value, consider the evolution of a typical computer system. When redesigning a computer that has its functional pieces interconnected the new artifact provides a single choice, the new system performs as a whole either better, or worse than its predecessor does. However, with the modularized version, the designer has the option to include each new module created for the next version, or leave it out, on a module by module basis. Furthermore, the modularization allows many experiments on the components most critical to overall system performance. The designer now has the *option* to pick the best outcome from many trials. For example, suppose the designers of a new computer system attempted a technically risky new technology for a CPU design, but it did not meet expectations, but rather had performance inferior to the previous version. The modular design allows using the old CPU, but also the *option* to include any improved components such as the display or memory systems. This approach is impossible with the interconnected version: the only *option* is to take or leave the entire new system. The modular design provides a portfolio of options rather than an option on a portfolio, which Black, Scholes, and Merton have shown has more value.

The value of this modularity is computed in [12]. Let V_1 be the value of a complex system built as a single module, and let V_j be the value of the same system with j modules. If we ignore the cost of modularity and make a few other assumtions (see [12]), then we get the value of dividing a complex system into j components is: $V_j = j^{1/2}V_1$. That is, the modularized system exceeds the value of the interconnected design by the square root of the number of modules.

4. A Model Applied to Networking Protocol Standards

Providing standardization options for different contingencies can help reduce the tremendous risk and associated costs of bad decisions regarding standardization. Below we present a generic structure that, if applied to standards, tends to create a standardization environment allowing the broadest range of experimenters to propose new standards. Our structure promotes end-to-end services, which tend to allow more experimentation than services within the network. Next, the market selects among the proposed standards, promulgating the standards most likely to be successful. Our methodology of structuring standards with a layered modularized architecture allowing applications with end-2-end like properties in regards to ease of innovation is quantifiable using the theory of real options. Lastly, we show the value of introducing standards in an evolutionary fashion by starting out with a simple version, and then growing the standard as the market evolves. Our theory is simple, is intuitive in nature, quantifiable, and empirically verifiable given the success of the Internet.

Our model depends on selection of standards from a market, but this market may be a set of vendors negotiating a standard. As discussed in Section 3 market selection, negotiated agreement, and hybrids between the two are selection mechanisms for standards. Even negotiated standards have an element of market selection in that the market can accept or reject what vendors agree to. OSI is a good example of this, vendors and governments negotiated the OSI specification within the guidance of the ISO, but users did not select it, instead choosing the Internet suite. In negotiated standards, one can view the organization responsible for the negotiations to be a market, selection the fittest technology in regards to vendors, but the user market has the right to reject the agreed to standard. Our theory works with both types of selection, it only requires multiple standards that can be picked by users, or agreed to by vendors.

4.1 Structure of a protocol stack

Figure 2 shows the structure of the Internet protocol suite at three points in its development. In (a), we see the TCP/IP Internet stack at its birth when it consisted of a single network/transport protocol (IP-TCP) and a few of several applications: ftp (a simple file transfer utility), and telnet (a remote login protocol). We date this at 1974 when Cerf and Kahn [19] published their first TCP/IP paper. Then, in (b) we see the next stage of evolution, the separating of IP and TCP into two layers (network and transport) and a new transport protocol (UDP) added for a different type of service to applications. This stage began with a hallway meeting in 1978 [20] about the merits of separating TCP and IP. Finally, in (c) is the Internet stack in December 1988 as specified in RFC 1083. This shows the growth of protocols of the Internet stack over a 14-year period.

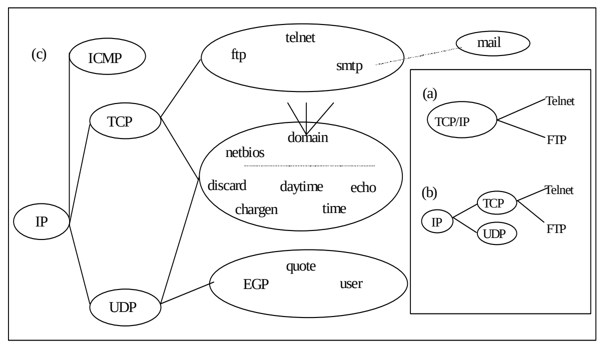

Figure 2 Internet suite structure - History

One of the most important attributes in this structure is what we call the "thin/thick" nature of the dependencies. Figure 2 (a) shows the most basic thin/thick structure, a single protocol (TCP/IP) providing a basic service usable by two very different application-layer protocols (ftp and telnet). Telnet and ftp have several things in common. First they need a reliable end-to-end data transport service provided by TCP/IP, and second, they are user-level applications, their development normally does not require changes to the network portion of the operating system. We define the fan-out as the number of protocols above using the lower layer service. In (b), the fan-out is two for both layers, and in (c), we see that the application-layer fan-out is growing much faster than for the layer below. The intuition behind this is that users want simple basic services that are broadly applicable to very diverse applications. Empirically this seems to be true with the Internet where TCP and UDP are the only standardized transport layer data transfer protocols in 20 years of Internet use.

This thin/thick structure falls out of the end-to-end argument [21], as this structure tends to push applications to the user layer with an end-to-end service model. The idea is to have a stupid network; the end systems provide the services and depend on the network only for data transport. This simplicity allows a broad range of researchers to experiment with application-layer protocols because they are simple to implement. The results of this structure can be amazing; just consider the creation of the web. Tim Berner-Lee was at CERN supporting computers for sub-particle atomic physics, not a network researcher. He invented both HTTP, HTML and the browser concept to better support his users' needs. The idea was so good that within ten years over 50% percent of total Internet traffic is based on web traffic [22].

4.2 General model assumptions

Since not all standards need to evolve quickly, or be modular in structure with a thin/thick structure, we present a set of conditions that when met by a technology implies that standardization of the technology will benefit from our methodology.

1. The conditions within which standardization occurs must have market uncertainty as described in Section 2.1
2. Technology used to implement the standard changes in predictable ways, but rapidly with short life cycles. The short life cycle of the technology requires flexibility and timeliness in creating standards.
3. The market for services enabled by the technology exists meaning that profit-seeking firms will have incentive to provide the service since customers are willing to pay for it, even in the absence of a standard. Without a market, market selection is not possible; the standard becomes anticipatory. While the general market exists, the particulars are uncertain. The particular feature set that will fit best with the market is unknown (i.e. email).

4.3 Hierarchical view of a protocol stack

The first step towards quantifying the value of a protocol standard is to define what to value, and what metric to use. At the top level is the value of the architecture of the stack describing the layers, modularization of protocols within each layer, and the dependency relationship between protocols in the layers. This protocol structure is similar to the design rules of modules discussed by Baldwin and Clark [12] and represents the conventions a protocol must adhere to if it is to function at a particular level in the stack[2]. One way to value the architecture of a protocol stack is in terms of how easy it is to augment the stack with new protocols, or change existing ones. Next in the hierarchy is the value of individual protocols within the suite. The value of a particular protocol is viewable in the context of how useful the protocol is for building services above it, the scalability, and ability to augment the protocol. Finally, the bottom layer is the value of a particular implementation of a standard or group or standards. Efficiency or speed of the implementation is a good metric for this bottom layer. The important point is that at each level there must be choices to make. These choices may be what type of structure the protocols should have within the stack, what feature set a particular protocol should have, or how well a particular implementation does compared to others in the context of performance and maintainability. While not suggesting these are the only measures, they are good examples. Our theory does not depend on the metric, but only on estimating the expected value and variance of the distribution describing the value.

The value of a standardized protocol suite, a single protocol, or the implementation of it is not deterministic due to uncertainty in users' preferences, and the unpredictability of engineers to create and enhance protocols in the prescribed manner. Thus, metrics such as market value or performance measures will be expected values, derived from the probability distribution of the outcome space. Real option theory provides a methodology to compute such expected values, with the value depending only on the variance and expected value of the distribution. In this case the variance of a distribution measures the market uncertainty.

4.4 Modularity

Our approach is similar to that used by Baldwin and Clark [12] showing the value of modular design over its interconnected cousin in computer design. The advantage of using modularity within each layer of a protocol stack is similar to the benefits gained by using modular design in computer systems. It allows keeping the best new module for a protocol (possibly picking the best outcome from many experiments) or keeping the old module, thus guaranteeing a higher expected value. To gain the most benefit from modularity there should be many choices for new modules. Architectures such as the end-2-end principal help in providing many choices because of the ease of experimentation.

Modularity of design and the thin/thick structure of protocol dependencies is particularly important to enhancing the value gained by augmentation of a protocol suite with standards for new protocols. Given the market uncertainty for network services, the more options a service provider has, the more likely it will meet demands of its users. By pushing applications to the user level with end-2-end applications, more experimentation is likely. There are several reasons for this. First, application-layer development is faster and less expensive than kernel work.

[2] Others have made arguments [Dclark90] [Bra] that functional layering is not the most efficient way to implement a protocol stack.

Next, the pool of talent with the skills to do application-layer coding is greater. Finally, the participants allowed to develop new services are much broader at the application level because users can innovate, and as Hippel [23] shows, users sometimes are best suited to solve their own problems. With thick/thin networks (i.e. like the telephone network) only those controlling the network can add new services. We do not provide a numerical example of this theory, but refer the reader to [12] to see how to accomplish this.

4.4.1 Cost of Modularity

The above arguments show that modularity is good, but it is hard, and can be expensive (as discussed in Section 3.1) which limits the granularity of modularity in a system. Defining modules that work together and are stable is very hard. Determining if different modules are compatible and will interoperate is expensive, the cost of testing modules and integrating them with the other protocols limits the number and complexity of modules existing at each level in the stack. This cost of modularity is the fixed cost of initially creating the module, plus the cost of testing each changed (or new) module, plus the cost of testing the integration of new modules with the entire system. Baldwin [12] shows how to factor this cost of modularity into the benefit of choice. We believe this cost depends on the layer within the stack the protocol belongs to and the number of other protocols that depend on this particular protocol. For example, in the Internet suite, all protocols above IP depend on IP, thus it is the most expensive protocol within the Internet to evolve. IP has changed little, and the acceptance of IPv6 has been slower than expected. Next are protocols at the transport layer like TCP and UDP that are relatively expensive to change (but do change, for example congestion control in TCP) since many different application-layer protocols use the same transport protocol. For example HTTP, FTP, and Telnet all use TCP for transport. Finally, the application-layer has the most protocols, and these are the least costly to change or add to (for example, HTTP). The dependence of the cost function upon the layer the protocol is in predicts the top half of the hourglass shape of the current Internet protocol suite, with IP being the common bearer service that glues everything together.

4.5 Staged Development

Another important attribute of the architecture of a protocol suite is its ability to evolve. Dyson says: "*we should not attempt to construct the Internet, but we should act like gardeners, providing a conducive environment for growth*" [24]. This suggests that protocol suites that evolve from a simple start will generally achieve a higher value than protocol suites (or individual protocols) that are initially complex. A protocol stack should start with as few protocols as needed to solve a current but focused problem. Email on the Internet is an example of a service that started out with simple protocols and evolved in complexity. It is far more successful than X.400, the OSI mail protocol that started with many more features[3]. Only after email had established itself did application protocols for transferring (SMTP) and accessing the email on the local email server (POP and IMAP) become standard Internet protocols. Furthermore, at first, only text-based e-mail was possible; only later did attachments (via MIME), allowing binary files to be sent as mail, become standardized. The first Internet mail specification (RFC 561) is 4 pages long compared to the current email specification (RFC 822) which is 46 pages long; the current MIME extensions (RFCs 2046 - 2049) comprise over 100 pages of specifications. This ability to evolve is essential to survival in uncertain environments. Unlike the unsuccessful X.400 protocol, Internet email protocols evolved into a set of standards that provided a feature set users wanted, and thus adopted.

One way to value an evolutionary style of enhancement to a protocol stack is to place this evolution in the context of a multi-staged investment. Similar to the example in Amran's *Real Option* [17] and related work by Kulatilaka [18], the evolution of a protocol stack is viewable as a series of staged investments. Each stage of development creates an option value by providing the choice of whether to continue evolving the stack, and how the protocol suite should change. Staging the investment required to develop a comprehensive set of protocols (or a complex single protocol) minimizes the risk of bad decisions in uncertainty.

Figure 3 shows the first two stages of a hypothetical evolution of protocols in the Internet. Stage one begins with market acceptance of a minimum set of protocols (i.e. TCP/IP). This stage has a single option, invest in a new transport layer protocol (UDP), or not. During stage one market uncertainty exists - will the market accept the new standard, or not? Stage two begins with four possible branches. At each stage of development we have a yes/no decision to make, and then roll the dice to see the outcome of our investment choice. This yes or no decision to continue with the protocol suite fits nicely with the binomial model. As the tree unfolds, we see all the various paths the evolution of the protocol can take. The example in [18] discusses how a dynamic programming algorithm back-solves this design tree to determine the optimal choice to make at each investment point. This strategy creates value by increasing the range of outcomes and providing relevant information that can be factored into

[3] This is particularly interesting given that X.400 will run on the Internet (via RFC 1006), but still did not become popular.

critical decisions. This provides a higher expected value than the single stage approach where there are only two possibilities: success of the full-blown protocol stack, or not. With standards it may be impossible to estimate the changes and choices that will arise in the evolution of protocols. However, this may be useful as a tool to perform a historic analysis of the particular evolutionary path a protocol has taken.

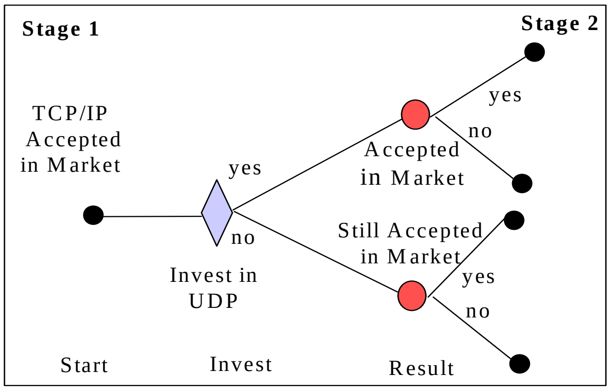

Figure 3 Multi-stage Protocol Introduction

This evolutionary approach gains from the modularity of each layer. Modularity of the protocol suite allows an additional option benefit at each decision point by permitting choice of the best of many proposed protocols to include in a stack. In effect, protocols developed with an evolutionary style and a layered modularized structure allow a double options benefit. That is, the option of whether to invest in a change, and if so what change to include next from the many choices.

This staged development model is applicable to studying the development of any particular protocol within the stack. The staged approach adds expected value to a protocol's evolution by allowing decisions about whether a protocol's development should continue, and if so, how to alter the protocol at each stage as more current information becomes available. One example of this is how TCP has evolved its sophisticated congestion control scheme. Initially, the congestion control was primitive, but once congestion existed, and was better understood, algorithms that are more effective became implemented [25].

This multi stage approach is not unlike how the IETF works with its tri-annual meetings. At each meeting working groups arrive at rough consensus about the solution to a technical problem, and then the solution advances along the standards track. Much of the work and changing status of Internet protocols occurs around the IETF meetings, with a flurry of activity before and after the conference.

4.6 High Level Empirical evidence

Empirical evidence from the Internet supports our model: it evolved as our model would predict. It is modular in design and layered such that it promotes end-2-end services, and the protocols started simple and evolved in complexity along with the users evolving expectations. The success of a protocol is its adoption by users. For example, as a whole the Internet stack is successful, while the OSI stack is not, but IS-IS is a successful OSI protocol, while Gopher is a failed Internet protocol.

The standard battle between the Internet protocol stack (promulgated by the IETF) and the OSI stack (standardized by the ISO) is a good example showing the success of technologies introduced as simple standards with few protocols and then allowed to grow in complexity. The Internet suite of protocols was introduced with only the basic building blocks, and these blocks had few options. In contrast, the OSI stack included protocols and options to the protocols to satisfy every possible need the designers could imagine. Initially introducing standards as "here is every thing you ever need" (as in the introduction of the OSI suite) requires over-standardization. Given the high probability of wrong predictions, the kitchen sink approach to standardization is very expensive. One example of this is the five-transport level protocols in the OSI stack, compared to the two in the Internet suite. The Internet protocol suite has shown the diversity of applications possible with the two extremes of the OSI suite: TP-0 is similar to UDP and TP-5, a TCP-like protocol. The OSI argument that different transport protocols are needed to efficiently handle networks of differing reliability turned out to be untrue [26]. Clark shows that the overhead a properly implemented heavy weight transport protocol like TCP occurs when packets arrive in order without data errors is only roughly 234 machine instructions. The market has spoken, OSI is dead for transport, and the Internet stack is the winner, yet many experts including Marshall Rose [27] believed differently, as late as the early 90s.

Another comparison to make is between the ITU's Frame Rely (introduced as a simple protocol with a successful evolution) WAN protocol and ATM. Frame Rely is an easy-to-understand WAN fast packet switching protocol that began as a short and simple specification from the ITU [I.122 (ANSI T1S1/88-224R)]. Frame Rely originally intended as a layer 2 Common Bearer service for ISDN, is seeing far more success in the WAN marketplace where it met a well-defined need - connecting LAN's over a wide area. Frame was versatile, it could be used to implement a private corporate network, or as an interface standard to connect to a public Frame Relay network. One competitor to Frame is ATM - the Holy Grail of low level networking protocols. ATM has it all, ultra fast bandwidth, fine-grained QoS, a complete solution from the desktop to the core. Unfortunately for ATM vendors, users did not need these advanced services and did not buy ATM.

Even with Frame's success, current use is far different than intended by the original providers of the service. Initially introduced as a medium bandwidth WAN service (in the .5 - 1.5 Mbps range) Frame Relay was unavailable as a lower speed (64K) service. Ironically, currently more than 50% of the Frame connections are 64K or less as a 1999 survey by the Frame Relay Forum shows. Furthermore, never-imagined applications like voice over Frame are evolving. Again, this shows that vendors cannot predict the demands of the users, or their willingness to pay for a service.

5. Real Options Applied to Protocol and Other IT Standards

Below are three general rules that help limit the risk of standardization of technology with market or other types of uncertainty based on our real options approach to standardization. The technology standardized should meet the general assumptions given in Section 4.2.

1. Standards should have a modularized (and possibly layered) architecture allowing the broadest range of experimentation in terms of number of experiments and groups able to contribute. The standardization process must allow market selection of the best outcomes.
2. A good way to introduce standards is in an evolutionary fashion; start out simple and build the complexity, thus allowing staged investment in creating and growing the standard.
3. Implementing a proposed standard is a good way to show it is possible. Furthermore, at least two independent implementations of the standard help show its clarity, completeness, and interoperability.

6. Conclusion

We have put forth a new paradigm for standardization under uncertainty using the theory of real options to quantify our results. Our work shows the value of a layered modularized protocol architecture initially introduced with a minimum of protocols and evolving in stages as market demands chart their chaotic path. Furthermore, the thin/thick structure of the stack pushes services to the end systems (end-2-end argument), allowing more experimentation by a broader range of participants. Introducing a protocol that solves a focused problem and then extending it in stages maximizes the expected value. The modular design of a protocol suite provides additional value by giving the designer a portfolio of options of many protocols from which to pick and choose. This is more valuable than a single option on a single complex protocol with many functions. Our argument extends to individual protocols within the stack, implying they should start out simple and be driven by market forces. Our quantitative prescriptive model is intuitively obvious and fits the empirical evidence of the Internet protocol suite and its early standardization process. The Internet Stack introduced in a staged fashion has maintained its basic thin/thick structure over its 20-year history.

7. Acknowledgments

Thanks to Carliss Baldwin, Barbara Feinberg, and Alan MacCormack - all from Harvard Business School for many great ideas, and brainstorming sessions. Carl Cargill from Sun MicroSystems for encouragement. Scott Brander for countless hours refining our ideas. Most of all H.T. Kung my advisor and Marco Iansiti (of HBS) for supporting this research, great ideas, and motivation, and editorial support from my wife Gretchen.

8. References

[1] Clark, Kim, The interaction of design hierarchies and market concepts in technological evolution, *Research Policy* 14 1985 p 235-251.

[2] Utterback, James, and Suarez, Fernando, Innovation, competition and industry structure, *Research Policy* 22 1993.

[3] Dosi, Giovanni, Technological paradigms and technological trajactories, *Research Policy* 11 p 147-162.

[4] Nelson, Richard, and Winter, Sidney, In search of useful theory of innovation, *Research Policy* 6, 1977 p 36-76.

[5] Anderson, Philip, and Tushman, Michael, Technological Discontinuities and Dominant Designs: A cyclical model of technological change. *Administrative Science Quarterly*, 35, 1990.

[6] Abernathy, William, and Clark, Kim, Innovation: Mapping the winds of creative destruction, *Research Policy* 14, 1985.

[7] Brown, Shona, and Eisenhardt, Kathleen, The Rate of Continuous Change: Linking Complexity Theory and Time-paced Evolution in Relentlessly Shifting Organizations, *Administrative Science Quarterly*, 42, 1997, p 1-37.

[8] Gould, Stephen, *Bully for Brontosaurus*, Norton & Company, NY 1992.

[9] Iansiti, Marco, Shooting the Rapids: managing product development in turbulent environments, *California Management Review* Vol 38, Num 1, Fall 1995.

[10] Rhoden, Desi, Real-Time On-Line Standards - The Ever Changing JEDEC and Other Standard Topics, Technology Standards & Standardization Processes, Their Impact on Innovation & International Competitiveness, *Proceedings of a Lecure Series Presented by US-Japan Technology Management Center*, Stanford University, June 1999.

[11] Cargill, Carl, *Open Systems Standardization, A Business Approach*, Prentice Hall, 1997.

[12] Baldwin, Carliss, and Clark, Kim (forthcoming*)*, *Design Rules: The Power of Modularity*. Cambridge, Ma: MIT Press 1999.

[13] Holland, John, *Adaptation in Natural and Artificial Systems*, Cambridge: MIT Press, 1992.

[14] Kauffman, Stuart, At Home In The Universe: *The Search for Laws of Self Organization and Complexity*, Oxford: University Press, 1993.

[15] Waldrop, Michell, *Complexity*, Simon & Schuster, 1992.

[16] Vercoulen, Frank, and Wegberg, Marc, Van, Standard Section Modes in Dynamic, Complex Industries: Creating Hybrids between Market Selection and Negotiated Selection of Standards, Working Paper Eindhoven University of Technology, School of Technology Management, Netherlands.

[17] Amram, Martha, and Kulatilaka, Nalin, *Real Options, Managing Strategic Investment in an Uncertain World*, Harvard Business School Press 1999.

[18] Kulatilaka, Nalin, and Balasubramanian, P., and Storck, John, *Managing Information Technology Investments: A Capability-based Real Options Approach,* School Of Management, Boston University.fd

[19] Cerf Vint, and Robert Kahn, A Protocol for Packet Network Interconnection, *IEEE Transactions of Communications*, 1974.

[20] Salus, Peter, *Casting the Net, From ARPANET to INTERNET and beyond…*, Addison-Wesley, 1995.

[21] Saltzer, Jerome, and David, Reed, and Clark, David, End-To-End Arguments in system design, *ACM Transactions in Computer Systems* 2, 4, Nov 1984 p 277-288.

[22] Berners-Lee, *Tim, Weaving the Web*, Harper San Francisco, 1999.

[23] Hippel, Eric, Economics of Product Development by Users: The Impact of "Sticky" Local Information, *Management Science*, Vol 44, No. 5, May 1998.

[24] Postrel, Virginia, *The Fruture and its Enemies,* Free Press Books, 1998.

[25] Jacobson, Van, Congestion Avoidance and Control. In *Proceedings, SIGCOMM'88 Workshop*, pages 314-329, ACM SIGCOMM, ACM Press, 1988.

[26] Clark, David, Jacobson, Van, and Romkey, John, and Salwen, Haward, An analysis of TCP processing overhead, *IEEE Communications Magazine*, June 1989.

[27] Rose, Marshall, The Open Book*, A Practical Perspective on OSI*, Prentice Hall, 1990.